

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
15 February 2001 (15.02.2001)

PCT

(10) International Publication Number
WO 01/11468 A1

(51) International Patent Classification: G06F 11/22, 11/20

MEYER, Douglas, B.: 3456 Old Cobble Court. San Diego, CA 92111 (US).

(21) International Application Number: PCT/US00/21743

(74) Agent: CONLEY, ROSE & TAYON, P.C.; B. Noel Kivlin, P.O. Box 398, Austin, TX 78767-0398 (US).

(22) International Filing Date: 8 August 2000 (08.08.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/371,059 9 August 1999 (09.08.1999) US


(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.

(71) Applicant: SUN MICROSYSTEMS, INC. [US/US]; 901 San Antonio Road, Palo Alto, CA 94303 (US).

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SI, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: DIAGNOSTIC CAGEDMODE FOR TESTING REDUNDANT SYSTEM CONTROLLERS



(57) Abstract: A multiprocessor system is disclosed that employs an apparatus and method for caging a redundant component to allow testing of the redundant component without interfering with normal system operation. In one embodiment the multiprocessor system includes at least two system controllers and a set of processing nodes interconnected by a network. The system controllers allocate and configure system resources, and the processing nodes each include a node interface that couple the nodes to the system controllers. The node interfaces can be individually and separately configured in a caged mode and an uncaged mode. In the uncaged mode, the node interface communicates information from either of the system controllers to other components in the processing node. In the caged mode, the node interface censors information from at least one of the system controllers. When all node interfaces censor information from a common system controller, this system controller is effectively "caged" and communications from this system controller are thereby prevented from reaching other node components. This allows the caged system controller along with all its associated

interconnections to be tested without interfering with normal operation of the system. Normal system configuration tasks are handled by the uncaged system controller. The uncaged system controller can instruct the node interfaces to uncage the caged system controller if the tests are successfully completed.

WO 01/11468 A1



Published:

With international search report.

Before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments.

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

DIAGNOSTIC CAGEMODE FOR TESTING REDUNDANT SYSTEM CONTROLLERS

BACKGROUND OF THE INVENTION

5 Field of the Invention

This invention relates to the field of multiprocessor computer systems with built-in redundancy, and more particularly, to systems and methods for testing redundant functional components during normal system operation.

Description of the Related Art

10 Multiprocessor computer systems include two or more processors which may be employed to perform computing tasks. A particular computing task may be performed upon one processor while other processors perform unrelated computing tasks. Alternatively, components of a particular computing task may be distributed among multiple processors to decrease the time required to perform the computing task as a whole. Generally speaking, a processor is a device that executes programmed instructions to produce desired output signals, often in response to
15 user-provided input data.

A popular architecture in commercial multiprocessor computer systems is the symmetric multiprocessor (SMP) architecture. Typically, an SMP computer system comprises multiple processors each connected through a cache hierarchy to a shared bus. Additionally connected to the shared bus is a memory, which is shared among the processors in the system. Access to any particular memory location within the memory occurs in a similar amount
20 of time as access to any other particular memory location. Since each location in the memory may be accessed in a uniform manner, this structure is often referred to as a uniform memory architecture (UMA).

Another architecture for multiprocessor computer systems is a distributed shared memory architecture. A distributed shared memory architecture includes multiple nodes that each include one or more processors and some local memory. The multiple nodes are coupled together by a network. The memory included within the multiple
25 nodes, when considered as a collective whole, forms the shared memory for the computer system.

Distributed shared memory systems are more scaleable than systems with a shared bus architecture. Since many of the processor accesses are completed within a node, nodes typically impose much lower bandwidth requirements upon the network than the same number of processors would impose on a shared bus. The nodes may operate at high clock frequency and bandwidth, accessing the network only as needed. Additional nodes may be
30 added to the network without affecting the local bandwidth of the nodes. Instead, only the network bandwidth is affected.

Because of their high performance, multiprocessor computer systems are used for many different types of mission-critical applications in the commercial marketplace. For these systems, downtime can have a dramatic and adverse impact on revenue. Thus system designs must meet the uptime demands of such mission critical
35 applications by providing computing platforms that are reliable, available for use when needed, and easy to diagnose and service.

One way to meet the uptime demands of these kinds of systems is to design in fault tolerance, redundancy, and reliability from the inception of the machine design. Reliability features incorporated in most multiprocessor computer systems include environmental monitoring, error correction code (ECC) data protection, and modular
40 subsystem design. More advanced fault tolerant multiprocessor systems also have several additional features, such as full hardware redundancy, fault tolerant power and cooling subsystems, automatic recovery after power outage, and advanced system monitoring tools.

For mission critical applications such as transaction processing, decision support systems, communications services, data warehousing, and file serving, no hardware failure in the system should halt processing and bring the whole system down. Ideally, any failure should be transparent to users of the computer system and quickly isolated by the system. The system administrator must be informed of the failure so remedial action can be taken to bring the computer system back up to 100% operational status. Preferably, the remedial action can be made without bringing the system down.

In many modern multiprocessor systems, fault tolerance is provided by identifying and shutting down faulty processors and assigning their tasks to other functional processors. However, faults are not limited to processors and may occur in other portions of the system such as, e.g., interconnection traces and connector pins. While these are easily tested when the system powers up, testing for faults while the system is running presents a much greater challenge. This may be a particularly crucial issue in systems that are "hot-swappable", i.e. systems that allow boards to be removed and replaced during normal operation so as to permit the system to be always available to users, even while the system is being repaired.

Further, some multiprocessor systems include a system controller, which is a dedicated processor or subsystem for configuring and allocating resources (processors and memory) among various tasks. Fault tolerance for these systems may be provided in the form of a "back-up" system controller. It is desirable for the primary and redundant system controllers to each have the ability to disable the other if the other is determined to be faulty. Further, it is desirable to be able to test either of the two subsystems during normal system operation without disrupting the normal system operation. This would be particularly useful for systems that allow the system controllers to be hot-swapped.

SUMMARY OF THE INVENTION

Accordingly, there is disclosed herein a multiprocessor system that employs an apparatus and method for caging a redundant component to allow testing of the redundant component without interfering with normal system operation. In one embodiment the multiprocessor system includes at least two system controllers and a set of processing nodes interconnected by a network. The system controllers allocate and configure system resources, and the processing nodes each include a node interface that couple the nodes to the system controllers. The node interfaces can be individually and separately configured in a caged mode and an uncaged mode. In the uncaged mode, the node interface communicates information from either of the system controllers to other components in the processing node. In the caged mode, the node interface censors information from at least one of the system controllers. When all node interfaces censor information from a common system controller, this system controller is effectively "caged" and communications from this system controller are thereby prevented from reaching other node components. This allows the caged system controller along with all its associated interconnections to be tested without interfering with normal operation of the system. Normal system configuration tasks are handled by the uncaged system controller. The uncaged system controller can instruct the node interfaces to uncage the caged system controller if the tests are successfully completed.

BRIEF DESCRIPTION OF THE DRAWINGS

A better understanding of the present invention can be obtained when the following detailed description of the preferred embodiment is considered in conjunction with the following drawings, in which:

Fig. 1 is a functional block diagram of a multiprocessor system; and

Fig. 2 is a functional block diagram of a processor node.

While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof are shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that the drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the intention is to cover all modifications, equivalents and alternatives falling within the spirit and scope of the present invention as defined by the appended claims.

DETAILED DESCRIPTION OF THE INVENTION

Turning now to the figures, Fig. 1 shows a block diagram of a multiprocessor system. The system includes a center plane 102 that interconnects N nodes (designated Node 0 through Node N-1) with a network bus (not shown). The network bus is preferably a crossbar network. The nodes preferably each include a node interface board 104 which accepts up to two boards one of which is designated as a "Slot 0 board" 106 while the other is designated as a "Slot 1 board" 108. Slot 0 boards are preferably multiprocessor boards that each include four processors, a memory module, and a system interface interconnected by a bus, and various support chips. Slot 1 boards are preferably I/O boards that interface to various peripherals such as serial and parallel ports, disk drives, modems, printers, etc.. In addition to the described types of Slot 0 and Slot 1 boards, other board types may be used, and the mix of the various board types among the various nodes is preferably alterable.

The system also includes at least two system controllers 110 which are preferably coupled to the center plane 102 by corresponding system controller support boards 112. The center plane 102 preferably provides busses from the support boards 112 to the nodes for maintenance, monitoring, and configuration of the nodes. The center plane 102 may also provide an arbitration bus 114 that allows the system controllers 110 to arbitrate for communication privileges to the nodes.

For a mission-critical system, it is necessary that the various components be hot-swappable so that defective components can be removed and replaced without bringing the system down. Accordingly, each of the node interface boards 104 and support boards 112, along with their dependent boards, can be removed and replaced while the system is operating. Since insertion is an event that has a relatively high failure probability, it is desirable to test the newly inserted components along with their physical interface to the system prior to trusting them with substantive tasks. The ensuing description focuses on testing of the system controllers 110 and support boards 112, but it is recognized that the nodes may be similarly tested.

Fig. 2 shows selected components common to each of the nodes. The node interface board 104 includes a system data interface chip 202, and each board 106, 108 includes a system data controller chip 204 that operates to configure and monitor various components on the board in accordance with information received from the system controller via the system data interface chip 202. The system data interface chip also operates to configure and monitor various components in the node interface 104 in accordance with communications received from the system controller. Both chips 202 and 204 are preferably able to parse address information and route communications from the system controller to the components indicated by the address information. The chips 202, 204 may additionally convert the communications into whatever form or bus protocol may be needed for the destination component to understand the message.

Referring concurrently to both Figs. 1 and 2, the system data interface chip 202 has a dedicated port for each of the system controllers 110, so that all communications with a given system controller are conducted via the associated port. The system data interface (SDI) chip 202 also includes some error detection and notification

circuitry. If the SDI chip 202 detects that a communication from a given system controller is corrupted, the SDI chip 202 can communicate an error notification to that system controller. However, if the SDI chip 202 is unable to determine the source of the error (e.g. when receiving conflicting communications from different system controllers) the SDI chip 202 may assert a system interrupt signal to alert the system controllers to the error event.

5 SDI chip 202 includes some status, configuration and test registers. The status registers may be read by the system controllers to determine error conditions, for example. One of the configuration registers includes "cage" mode bits that can be asserted and de-asserted only by an "uncaged" system controller. An uncaged system controller may put one or all of its interfaces into a cage mode, but an uncaged system controller will be required to put them back into an uncaged mode. It is noted that in situations where both node interfaces are caged, or a caged
10 system controller can not respond to a command to exit the cage, either system controller (whether caged or not) can initiate a bus reset that will force the node interface back to an uncaged mode.

Either of the system controllers can be caged by assertion of an associated cage mode bit. The assertion of cage mode bits may be accomplished by one of the system controllers writing an individual caging message to each of the nodes. The SDI chips 202 in each of the nodes interpret the caging message and assert the cage mode bit for
15 the designated system controller. The system controller designated in the caging message to a node interface is hereafter referred to as a caged system controller for that node interface. Conversely, a system controller for which the cage mode bits in a node interface are not asserted is hereafter referred to as an uncaged system controller for that node interface.

Either of the system controllers can have one or more of its interfaces caged by writing a cage enable
20 message to the pertinent node interfaces. If all node interfaces have the same system controller interface caged, then the system controller is said to be completely caged. If not all not all node interfaces have the same system controller interface caged, then the system controller is incompletely caged, and it is permitted to communicate with interfaces for which it is uncaged.

Assertion of a cage mode bit causes the SDI chip 202 to censor any communications received from the
25 caged system controller. The SDI chip 202 may communicate responses to the caged system controller such as, e.g. error notification for corrupted communications. The SDI chip 202 may also operate on communications from the caged system controller, e.g. storing values in the test registers. However, the SDI chip 202 does not transmit any messages to other downstream components in response to communications received from the caged system controller. This includes configuration messages for the boards 106, 108, as well as messages for other components
30 on node interface board 104. The SDI chip also suppresses interrupts triggered by communications from the caged system controller, such as a protocol error interrupt that would normally be caused by a message from the caged system controller that conflicts with a message received from the other system controller.

When the multiprocessor computer system is first powered up, the primary system controller runs through a Power-On-Self-Test (POST) which tests all of the components on the system controller and then tests all of the
35 interconnections to other boards/components in the multiprocessor system. Since no user applications are active, serious or fatal errors will not cause service interruptions. However, if the multiprocessor system is executing user applications and a secondary system controller needs to be tested, then the caging mode may be employed to test the secondary system controller while the primary system controller continues providing services to the hardware and mastering all maintenance buses required for testing. The caging mode will prevent the system controller under
40 test from inadvertently destroying state information in the active hardware or causing an error which would not be

isolated and would be reported as an system error from the component under test. Such actions would probably bring down the multiprocessor system.

Referring to Fig. 1, a newly inserted system controller support board 112 with attached system controller 110 is caged by placing all of its node interfaces into caging mode. This can be done by the newly inserted system controller or by the resident system controller. A test process executing on the caged system controller is then able to verify the functionality of not only the on-board components, but also the components on the support board 112, the center plane 102, and portions of the SDI chip 202. It is noted that the interconnections between the system controller 110, the support board 112, the center plane 102, and the node interface 104 are also verified by the test process. The uncaged system controller can check for successful completion of the test process, e.g. by reading status registers in the caged system controller and/or status registers in the SDI chips 202, and broadcast an uncaging message to the SDI chips 202 if the test process is determined to have completed successfully.

In addition to testing itself, the caged system controller is able to test off-board interconnects without concern of interfering with running software applications and the primary system controller. Without this ability, the caged system controller could not detect faulty interconnections with the rest of the system. If untested interconnections to the inserted system controller were faulty, this would not be known until after the primary system controller had failed. At that point the faults would appear and the system would probably crash. The detection of interconnect faults before failure of the primary system controller allows time for notification and for remedial action.

It is noted that discussion has centered on testing of a redundant system controller by placing all node interfaces into the caging mode. However, the described embodiment also allows node interfaces to be separately and individually placed into the caging mode. This allows testing of individual bus connections while the system controller is able to maintain its duties elsewhere.

One embodiment of the invention has been generally described above. The following discussion describes various details of one particular preferred implementation for explanatory purposes. However, the invention is not so limited.

The invention may be employed in a next generation, UltraSPARC III based, high-end enterprise server system. The system controllers may be single processor based subsystems which provide many global resources to all of the multiprocessor hardware. The system controllers may employ a variety of buses to obtain complete access to all hardware. Preferably, more than one system controller is present at any given time one but only one is active. The second system controller preferably waits in a stand-by mode in case the primary system controller experiences a hardware failure.

The system controller interconnections to other hardware occur through various buses such as I2C (Inter-Integrated Circuit), JTAG (Joint Test Activity Group), and Console Bus. In a normal (uncaged) operating mode, the node interfaces multiplex together both system controller's Console Buses and provide no hardware isolation. Thus, all boards and components in the system see all transactions emanating from either system controller. In caging mode, the node interfaces isolate a system controller and its Console Bus interconnections to the various hardware boards to prevent faults and protocol errors from propagating. This allows the system to properly test a system controller board while the system is running user application programs without causing a complete system crash.

The center plane may be a 16x16 crossbar interconnection network such as Sun Microsystems' Inc. Gigaplane-XB. This center plane contains two symmetrical sides which can each mount up to eight system boards, a support board and a system controller board. The system boards reside on node interface boards that connect to

the center plane through I2C bus and Console Bus. The I2C bus is a serial data bus developed by Philips Corporation consisting of a two line interface. One line consists of a data pin for input and output functions and the other line is a clock for reference and control.

Console Bus is a bus developed by Sun Microsystems Inc. and is used by the system controller as a pathway for status and control of all system functions. The System Data Interface (SDI) chip contains a console bus interface used by the system controller as a pathway for status monitoring and configuration control of all system functions. The console bus is the primary system control/diagnostics bus and is required to operate correctly at all times while the system is operational. Dual console bus interfaces, one to each system controller, are provided for redundancy.

Because of its critical importance, the SDI also contains a console bus cage mechanism to facilitate diagnostic testing of one of the two console bus interfaces while the other console bus interface is actively being used by the system for monitoring and configuration functions. Additionally, both interfaces of an SDI chip may be caged and tested independently if the situation requires (e.g. when a new node is inserted into a working system). The console bus cage operates to ensure that any event (correct or erroneous) that occurs while accessing a caged console bus has no impact on the normal functioning of the rest of the system, and specifically not on the other console bus operations. If a system controller after being caged and tested is not functioning correctly, the uncaged system controller can access SDI status registers that contain diagnostic identification information to determine the nature of the errors.

During normal operation, the uncaged Console Bus interface in the SDI chip handles any address translations required from the node interface board to either of the resident Slot 0 and Slot 1 boards. In this mode, a single state machine may be shared between the SDI console bus ports. In the caged mode, a second state machine may be used to handle transactions with the caged system controller. The transition from uncaged mode to caged mode can occur at any time. However, to avoid protocol errors, the transition from the caged mode to uncaged mode can occur only when the uncaged mode state machine is quiescent. A cage control register includes bits for indicating the activity of the state machines and the cage mode bits for caging the system controllers.

All accesses from a caged system controller are examined to determine if they are within the allowed range of addresses for status and test registers. Accesses outside this range are responded to with an illegal-access-error acknowledgement, and the access is suppressed. Error notifications may be posted in an error (status) register, but no interrupts are caused by caged transactions.

An example of a Slot 0 board which may be employed in the system is multiprocessor system board that holds four Sun Microsystems UltraSPARC III microprocessors with supporting level two cache. An example of a Slot 1 board which may be employed in the system is an I/O board with multiple PCI interfaces with slots for networking and I/O adapters. PCI bus is a standard bus used in computer systems to communicate data, instructions and control information between logic circuits.

The system controller support boards connect the system controllers to the center plane through the Console Bus and the I2C bus. The system controller support boards are repeaters, that is, they amplify and multiply the output signals from the system controller board and send them to the center plane for output to the node interface boards. The system controller board contains some system-level logic that includes a system clock generator, temperature and airflow monitoring circuitry, and a PCI interface to a computer system which handles diagnostics, boot, shutdown and environmental monitoring. The multiprocessor computer system requires only one

system control board for proper operation. However, for higher levels of system availability, a second optional system control board may be installed.

The computer system included in the system controller board includes an UltraSparc Ili microprocessor and various programmable read only memories (PROM) containing software for configuration and testing of the hardware in the multiprocessor computer system. The system level logic converts PCI signals into I2C and Console Bus and, after amplification and multiplication in the support board, sends these signals to the center plane. The system-level logic also controls JTAG scan chains which connect through the center plane and all hardware boards in the multiprocessor computer system. JTAG test access ports are present throughout the center plane and various boards of the multiprocessor computer system and allow for greater visibility and verification of system boards when the system controller performs the POST.

During operation of the multiprocessor computer system, a replacement microprocessor board or I/O board after it has been inserted must be attached electrically to the remainder of the hardware. The component must be isolated from the other hardware present in the computer system and tested prior to and during attachment. Finally, the hardware component must be incorporated logically into the running multiprocessor computer system to run the operating system and execute application programs for users.

In one preferred embodiment, the replacement microprocessor or I/O board becomes a part of a dynamic system domain after it has been inserted into the center plane of the multiprocessor computer system. Dynamic system domains are software partitions that permit the multiprocessor computer system to be dynamically subdivided into multiple computers. A dynamic system domain may consist of one or more system boards. Each domain is a separate shared-memory SMP system that runs its own local copy of a multiprocessor operating system such as Sun Microsystems Solaris and has its own disk storage and network connections. Because individual system domains are logically isolated from other system domains, hardware and software errors are confined to the domain in which they occur and do not affect the rest of the system. After a system administrator requests a particular domain configuration, the system controller configures the various microprocessor and I/O boards into dynamic system domains in the multiprocessor computer system.

Modifications to the hardware makeup of a domain may be required while the multiprocessor computer system is in operation. In order to facilitate run time changes in dynamic system domain configuration, system administrators should be able to dynamically switch system boards between domains or remove them from active domains for testing, upgrades, or servicing. Ideally after testing or service hardware boards should be easily reintroduced into one of the active domains without interrupting system operation. Each system domain is administered through the system controller which services all the domains. The system controllers may interface to a SPARC workstation or equivalent computer system board that runs a standard operating system such as Microsoft Windows NT or Microsoft Windows 98, Sun Microsystems Solaris, IBM AIX, Hewlett Packard UX or some similar equivalent and a suite of diagnostic and management programs. The external computer system may be connected via a network interface card such as Ethernet to the system controller located in the multiprocessor computer system. The microprocessor in the system controller board interprets the network interface card (e.g. TCP/IP Ethernet) traffic and converts it to encoded control information.

Numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.

WHAT IS CLAIMED IS:

1. A multiprocessor computer system that comprises:

a first and second system controllers each configurable to allocate and configure system resources; and

a plurality of processing nodes interconnected by a communications link, wherein each of the processing nodes is further coupled to both the first and second system controllers, wherein each of the processing nodes includes:

a node interface configurable between a caged mode and an uncaged mode, wherein the node interface communicates information from either of the system controllers to other node components only when in the uncaged mode, and wherein the node interface censors information from a selected one of the system controllers when in the caged mode.

2. The multiprocessor computer system of claim 1, wherein each of the node interfaces is switched from the uncaged mode to the caged mode to cage the selected system controller.

3. The multiprocessor computer system of claim 1, wherein the node interface is further configurable between a connection mode and an isolation mode, wherein the node interface conveys information from other node components to the communications link when in connection mode, and wherein the node interface suppresses information from other node components when in the isolation mode.

4. The multiprocessor system of claim 1, wherein the node interfaces each include one or more registers configurable to store information from the selected system controller and configurable to communicate the stored information to a requesting one of the system controllers.

5. The multiprocessor system of claim 1, wherein the communications link is provided by a center plane that interconnects circuit boards which form the processing nodes.

6. The multiprocessor system of claim 5, wherein one or more of the processing nodes include:

a processor board that has multiple processors and a memory module configured in a shared bus architecture, wherein the multiple processors communicate with the center plane via a bus bridge included on the processor board.

7. The multiprocessor system of claim 6, wherein the one or more processing nodes further include:

a node interface board coupled between the processor board and the center plane, wherein the node interface board includes the node interface.

8. The multiprocessor system of claim 7, wherein the one or more processing nodes further include:

an I/O board coupled to the node interface board.

9. The multiprocessor system of claim 7, wherein the system controllers are each coupled to the center plane by a system controller support board, wherein each of the system controllers are further coupled to each other for arbitration via the center plane and the system controller support boards.

5 10. The multiprocessor system of claim 9, wherein the first system controller is configured to instruct the node interfaces to enter the caged mode to censor information from the second system controller before the second system controller performs testing of associated system controller support board, center plane and node interface board connections.

10 11. The multiprocessor system of claim 1, wherein the first system controller is configured to instruct the node interfaces to enter the caged mode to censor information from the second system controller if the first system controller determines that the second system controller is faulty.

12. The multiprocessor system of claim 1, wherein when in the uncaged mode the node interface is configured to
15 assert a system interrupt upon detecting an error in information received from either one of the system controllers, and wherein when in the caged mode the node interface does not assert the system interrupt in response to errors in information received from the caged system controller.

13. A method for verifying component functionality during system operation, wherein the method comprises:

20 caging a redundant portion of the system by placing interface elements between the redundant system portion and a remaining portion of the system into a cage mode, wherein the interface elements are configured to block communications from the redundant component when in the cage mode, and are further configured to convey communications from the redundant component to the remaining portion of the system when in an uncaged mode; and

25 initiating a test process on the redundant system portion while normal processes continue operating on the remaining system portion.

14. The method of claim 13, wherein the redundant portion of the system includes:

30 a system controller configurable to allocate and configure system resources;
a system controller interface board coupled between the system controller and a center plane;
bus lines on the center plane that are associated with the system controller interface; and
node interface board connectors coupled to said bus lines.

15. The method of claim 14, wherein the interface elements are node interface boards coupled between the center
35 plane and respective processor boards.

16. The method of claim 13, further comprising:

40 detecting insertion of a component of the redundant system portion before caging the redundant system portion; and
uncaging the redundant system portion if the test process indicates that the redundant portion is functional.

17. A redundant system that comprises:

a resource allocation means for allocating and configuring system resources;

a plurality of processing means for accomplishing assigned tasks, wherein said plurality of processing means is interconnected by communications means, wherein each of said processing means is further coupled to the resource allocation means, and wherein each processing means includes:

an interface means for configuring components of the processing means in response to signals from the resource allocation means, wherein the interface means is configurable between a caged mode and an uncaged mode, wherein the interface means communicates signals from the resource allocation means to said components of the processing means only when configured in the uncaged mode, and wherein the interface means censors information from the resource means when in the caged mode.

18. The redundant system of claim 17, wherein the interface means each include one or more registers configurable to store information from the resource allocation means, and configurable to communicate the stored information to a requesting resource allocation means.





FIG. 1

**FIG. 2**

INTERNATIONAL SEARCH REPORT

Inte. .onal Application No

PCT/US 00/21743

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F11/22 G06F11/20

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 G06F H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

PAJ, EPO-Internal, WPI Data, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 4 342 084 A (SAGER GORDON S ET AL) 27 July 1982 (1982-07-27)	1
X	abstract	13
Y	US 4 684 885 A (CHAPMAN ET AL.) 4 August 1987 (1987-08-04)	1
X	the whole document	13, 17
A	US 4 914 572 A (BITZINGER RUDOLF ET AL) 3 April 1990 (1990-04-03)	1-18
A	US 5 765 034 A (RECIO RENATO JOHN) 9 June 1998 (1998-06-09)	1-18
	the whole document	
	-/--	

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

G document member of the same patent family

Date of the actual completion of the international search

15 December 2000

Date of mailing of the international search report

27/12/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Absalom, R

INTERNATIONAL SEARCH REPORT

Inte. onal Application No
PCT/US 00/21743

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 623 643 A (JANSSEN PAUL E ET AL) 22 April 1997 (1997-04-22) the whole document ---	1-18
A	US 5 416 921 A (SHEPARD WILLIAM C ET AL) 16 May 1995 (1995-05-16) the whole document -----	1-18

INTERNATIONAL SEARCH REPORT

Information on patent family members

Inte. .onal Application No

PCT/US 00/21743

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 4342084 A	27-07-1982	DE 3176048 D EP 0045836 A JP 1421165 C JP 57037800 A JP 62029827 B	30-04-1987 17-02-1982 29-01-1988 02-03-1982 29-06-1987
US 4684885 A	04-08-1987	BE 905688 A CA 1263759 A IT 1213507 B	02-03-1987 05-12-1989 20-12-1989
US 4914572 A	03-04-1990	AT 71788 T BR 8701106 A CN 87101838 A,B DE 3775946 A DE 3775946 D DK 124687 A EP 0236803 A FI 871058 A,B, GR 3004339 T PT 84444 A,B	15-02-1992 29-12-1987 02-12-1987 27-02-1992 27-02-1992 13-09-1987 16-09-1987 13-09-1987 31-03-1993 01-04-1987
US 5765034 A	09-06-1998	NONE	
US 5623643 A	22-04-1997	CA 2107755 A,C EP 0598526 A JP 6214906 A	17-05-1994 25-05-1994 05-08-1994
US 5416921 A	16-05-1995	NONE	